

Quantile Reinforcement Learning

Hugo Gilbert

*Sorbonnes Universités
UPMC Univ Paris 06
CNRS, LIP6 UMR 7606
Paris, France*

HUGO.GILBERT@LIP6.FR

Paul Weng

*School of Electronics and Information Technology
SYSU-CMU Joint Institute of Engineering
SYSU-CMU Shunde Joint Research Institute
Guangzhou, 510006, PR China*

PAWENG@CMU.EDU

Abstract

In reinforcement learning, the standard criterion to evaluate policies in a state is the expectation of (discounted) sum of rewards. However, this criterion may not always be suitable, we consider an alternative criterion based on the notion of quantiles. In the case of episodic reinforcement learning problems, we propose an algorithm based on stochastic approximation with two timescales. We evaluate our proposition on a simple model of the TV show, Who wants to be a millionaire.

Keywords: Reinforcement learning, Quantile, Ordinal Decision Model, Two-Timescale Stochastic Approximation

1. Introduction

Markov decision process and reinforcement learning are powerful frameworks for building autonomous agents (physical or virtual), which are systems that make decisions without human supervision in order to perform a given task. Examples of such systems abound: expert backgammon player (Tesauro, 1995), dialogue systems (Zhang et al., 2001), acrobatic helicopter flight (Abbeel et al., 2010) or human-level video game player (Mnih et al., 2015). However, the standard framework assumes that rewards are numeric, scalar and additive and that policies are evaluated with the expectation criterion. In practice, it may happen that such numerical rewards are not available, for instance, when the agent interacts with a human who generally gives ordinal feedback (*e.g.*, “excellent”, “good”, “bad” and so on). Besides, even when this numerical information is available, one may want to optimize a criterion different than the expectation, for instance in one-shot decision-making.

Several works considered the case where preferences are qualitative. Markov decision processes with ordinal reward have been investigated (Weng, 2012, 2011) and different ordinal decision criteria have been proposed in that context. More generally, preference-based reinforcement learning (Akrouf et al., 2012; Fürnkranz et al., 2012; Busa-Fekete et al., 2013, 2014) has been proposed to tackle situations where the only available preferential information concerns pairwise comparisons of histories.

In this paper, we propose to search for a policy that optimizes a quantile instead of the expectation. Intuitively, the τ -quantile of a distribution is the value q such that the probability of getting a value lower than q is τ (and therefore the probability of getting a value greater than q is $1 - \tau$). The median is an example of quantile where $\tau = 0.5$. Interestingly, in order to use this criterion only an order over valuations is needed.

The quantile criterion is extensively used as a decision criterion in many domains. In finance, it is a risk measure and is known as Value-at-Risk (Jorion, 2006). For its cloud computing services, Amazon reports (DeCandia et al., 2007) that they optimize the 99.9%-quantile¹. In fact, decisions in the web industry are often made based on quantiles (Wolski and Brevik, 2014; DeCandia et al., 2007). More generally, in the service industry, because of skewed distributions (Benoit and Van den Poel, 2009), one generally does not want that customers are satisfied on average, but rather that most customers (*e.g.*, 99% of them) to be as satisfied as possible.

The use of the quantile criterion can be explained by the nice properties it enjoys:

- preferences and uncertainty can be valued on scales that are not commensurable,
- preferences over actions or trajectories can be expressed on a purely ordinal scale,
- preferences over policies are more robust than with the standard criterion of maximizing the expectation of cumulated rewards.

The contributions of this paper are as follows. To the best of our knowledge, we are the first to propose an RL algorithm to learn a policy optimal for the quantile criterion. This algorithm is based on stochastic approximation with two timescales. We present an empirical evaluation of our proposition on a version of Who wants to be a millionaire.

The paper is organized as follows. Section 2 presents the related work. Section 3 recalls the necessary background for presenting our approach. Section 4 states the problem we solve and introduce our algorithm, Quantile Q-learning. Section 5 presents some experimental results. Finally, we conclude in Section 6.

2. Related Work

A great deal of research on MDPs (Boussard et al., 2010) considered decision criteria different to the standard ones (*i.e.*, expected discounted sum of rewards, expected total rewards or expected average rewards). For instance, in the operations research community, White (1987) notably considered different cases where preferences over policies only depend on sums of rewards: Expected Utility (EU), probabilistic constraints and mean-variance formulations. In this context, he showed the sufficiency of working in a state space augmented with the sum of rewards obtained so far. Filar et al. (1989) investigated decision criteria that are variance-penalized versions of the standard ones. They formulated the obtained optimization problem as a non-linear program. Yu et al. (1998) optimized the probability that the total reward becomes higher than a certain threshold.

Additionally, in the artificial intelligence community, Liu and Koenig (2005, 2006) also investigated the use of EU as a decision criterion in MDPs. To optimize it, they proposed a

1. Or 0.01%-quantile, depending on whether the problem is expressed in terms of costs or rewards.

functional variation of Value Iteration. In the continuation of this work, [Gilbert et al. \(2015\)](#) investigated the use of Skew-Symmetric Bilinear (SSB) utility ([Fishburn, 1981](#)) functions — a generalization of EU that enables intransitive behaviors and violation of the independence axiom — as decision criteria in finite-horizon MDPs. Interestingly, SSB also encompasses probabilistic dominance, a decision criterion that is employed in preference-based sequential decision-making ([Akrouf et al., 2012](#); [Förnkrantz et al., 2012](#); [Busa-Fekete et al., 2013, 2014](#)).

In theoretical computer science, sophisticated decision criteria have also been studied in MDPs. For instance, [Gimbert \(2007\)](#) proved that many decision criteria based on expectation (of limsup, parity... of rewards) admit a stationary deterministic optimal policy. [Bruyère et al. \(2014\)](#) considered sophisticated preferences over policies, which amounts to searching for policies that maximize the standard criterion while ensuring an expected sum of rewards higher than a threshold with probability higher than a fixed value. This work has also been extended to the multiobjective setting ([Randour et al., 2014](#)).

Recent work in Markov decision process and reinforcement learning considered conditional Value-at-risk (CVaR), a criterion related to quantile, as a risk measure. [Bäuerle and Ott \(2011\)](#) proved the existence of deterministic wealth-Markovian policies optimal with respect to CVaR. [Chow and Ghavamzadeh \(2014\)](#) proposed gradient-based algorithms for CVaR optimization. In contrast, [Borkar and Jain \(2014\)](#) used CVaR in constraints instead of as objective function.

Closer to our work, several quantile-based decision models have been investigated in different contexts. In uncertain MDPs where the parameters of the transition and reward functions are imprecisely known, [Delage and Mannor \(2007\)](#) presented and investigated a quantile-like criterion to capture the trade-off between optimistic and pessimistic viewpoints on an uncertain MDP. The quantile criterion they use is different to ours as it takes into account the uncertainty present in the parameters of the MDP.

In MDPs with ordinal rewards ([Weng, 2011, 2012](#); [Filar, 1983](#)), quantile-based decision models were proposed to compute policies that maximize a quantile using linear programming. While quantiles in those works are defined on distributions over ordinal rewards, quantiles in this paper are defined on distributions over histories.

More recently, in the machine learning community, quantile-based criteria have been proposed in the multi-armed bandit (MAB) setting, a special case of reinforcement learning. [Yu and Nikolova \(2013\)](#) proposed an algorithm in the pure exploration setting for different risk measures, including Value-at-Risk. [Carpentier and Valko \(2014\)](#) studied the problem of identifying arms with extreme payoffs, a particular case of quantiles. Finally, [Szörényi et al. \(2015\)](#) investigated MAB problems where a quantile is optimized instead of the mean.

The algorithm we propose is based on stochastic approximation with two timescales, a technique proposed by [Borkar \(1997, 2008\)](#). This method has recently been exploited in achievability problems ([Blackwell, 1956](#)) in the context of multiobjective MDPs ([Kalathil et al., 2014](#)) and for learning SSB-optimal policies ([Gilbert et al., 2016](#)).

3. Background

We provide in this section the background information necessary to present our algorithm to learn a policy optimal for the quantile criterion.

3.1. Markov Decision Process

Markov Decision Processes (MDPs) offer a powerful formalism to model and solve sequential decision-making problems (Puterman, 1994). A finite horizon MDP is formally defined as a tuple $\mathcal{M}_T = (S, A, \mathbf{P}, \mathbf{R}, s_0)$ where:

- T is a finite horizon,
- S is a finite set of states,
- A is a finite set of actions,
- $\mathbf{P} : S \times A \times S \rightarrow \mathbb{R}$ is a transition function with $\mathcal{P}(s, a, s')$ being the probability of reaching state s' when action a is performed in state s ,
- $\mathbf{R} : S \times A \rightarrow \mathbb{R}$ is a bounded reward function and
- $s_0 \in S$ is an initial state.

In this model, starting from initial state s_0 , an agent chooses at every time step t an action a_t to perform in her current state s_t , which she can observe. This action results in a new state $s_{t+1} \in S$ according to probability distribution $\mathbf{P}(s_t, a_t, \cdot)$, and a reward signal $\mathbf{R}(s_t, a_t)$, which penalizes or reinforces the choice of this action.

We will call t -*history* h_t a succession of t state-action pairs starting from state s_0 (e.g., $h_t = (s_0, a_1, s_1, \dots, s_{t-1}, a_{t-1}, s_t)$). The action choices of the agent is guided by a policy. More formally, a *policy* π at an horizon T is a sequence of T decision rules $(\delta_T, \dots, \delta_1)$. *Decision rules* prescribe which action the agent should perform at a given time step. They can be *Markovian* if they only depend on the current state. Besides, a decision rule is either *deterministic* if it always selects the same action in a given situation or *randomized* if it prescribes a probability distribution over possible actions. Consequently, a policy can be *Markovian*, *deterministic* or *randomized* according to the type of its decision rules. Lastly, a policy is *stationary* if it applies the same decision rule at every time step, i.e., $\pi = (\delta, \delta, \dots)$.

Policies can be compared with respect to different decision criteria. The usual criterion is the *expected (discounted) sum of rewards*, for which an optimal deterministic Markovian policy is known to exist for any horizon T . This criterion is defined as follows. First, the value of a history $h_t = (s_0, a_1, s_1, \dots, s_{t-1}, a_t, s_t)$ is described as the (possibly discounted) sum of rewards obtained along it, i.e.,

$$r(h_t) = \sum_{i=1}^t \gamma^{i-1} \mathbf{R}(s_{i-1}, a_i)$$

where $\gamma \in [0, 1]$ is a discount factor. Then, the value of a policy $\pi = (\delta_T, \dots, \delta_1)$ in a state s is set to be the expected value of the histories that can be generated by π from s . This value, given by the *value function* $v_T^\pi : S \rightarrow \mathbb{R}$ can be computed iteratively as follows:

$$\begin{aligned} v_0^\pi(s) &= 0 \\ v_t^\pi(s) &= \mathbf{R}(s, \delta_t(s)) + \gamma \sum_{s' \in S} \mathbf{P}(s, \delta_t(s), s') v_{t-1}^\pi(s') \end{aligned} \tag{1}$$

The value $v_t^\pi(s)$ is the expectation of cumulated rewards obtained by the agent if she performs action $\delta_t(s)$ in state s at time-step t and continues to follow policy π thereafter. The higher the values of $v_t^\pi(s)$ are, the better. Therefore, value functions induce a preference relation \succsim_π over policies in the following way:

$$\pi \succsim_\pi \pi' \Leftrightarrow \forall s \in S, v_T^\pi(s) \geq v_T^{\pi'}(s)$$

A solution to an MDP is a policy, called *optimal policy*, that ranks the highest with respect to \succsim_π . Such a policy can be found by solving the following equations, which yields the value function of an optimal policy:

$$\begin{aligned} v_0^*(s) &= 0 \\ v_t^*(s) &= \max_{a \in A} \mathbf{R}(s, a) + \gamma \sum_{s' \in S} \mathbf{P}(s, a, s') v_{t-1}^*(s') \end{aligned}$$

3.2. Reinforcement Learning

In the reinforcement learning (RL) setting, the assumption of the knowledge of the environment is relaxed: both dynamics through the transition function and preferences via the reward function are not known anymore. While interacting with its environment, an RL agent tries to learn a good policy by trial and error.

To make finite horizon MDPs learnable, we assume the decision process is repeated infinitely many times. That is, when horizon T is reached, we assume that the agent automatically returns to the initial state and the problem starts over.

A simple algorithm to solve such an RL problem is the Q-learning algorithm (see Algorithm 1), which estimates the Q-function:

$$Q_t(s, a) = \mathbf{R}(s, a) + \gamma \sum_{s' \in S} \mathbf{P}(s, a, s') V_{t-1}(s')$$

And obviously, we have: $V_t(s) = \max_{a \in A} Q_t(s, a)$.

In Algorithm 1, Line 1 generally depends on the $Q_{t-1}(s, \cdot)$ and possibly on iteration n . A simple strategy to perform this choice is called *ε -greedy* where the best action dictated by $Q_{t-1}(s, \cdot)$ is chosen with probability $1 - \varepsilon$ (with ε a small positive value) or a random action is chosen otherwise. A schedule can be defined so that parameter ε tends to zero as n tends to infinity. Besides, $\alpha_n(s, a) \in (0, 1)$ on Line 2 is a learning rate. In the general case, it depends on iteration n , state s and action a , although in practice it is often chosen as a constant.

3.3. Limits of standard criteria

The standard decision criteria used in MDPs, which are based on expectation, may not be reasonable in some situations. Firstly, unfortunately, in many cases, the reward function \mathbf{R} is not known. In those cases, one can try to recover the reward function from a human expert (Ng and Russell, 2000; Regan and Boutilier, 2009; Weng and Zanuttini, 2013). However, even for an expert user, the elicitation of the reward function can reveal burdensome. In inverse reinforcement learning (Ng and Russell, 2000), the expert is assumed to know an

Data: $\mathcal{M}_T = (S, A, G, \mathbf{P}, \mathbf{R}, s_0)$ MDP
Result: Q
begin
 $Q_0(s, a) \leftarrow 0, \forall (s, a) \in S \times A$
 $s \leftarrow s_0$
 $t \leftarrow 1$
for $n = 1$ **to** N **do**
1 $a \leftarrow$ choose action
 $r, s' \leftarrow$ perform action a in state s
2 $Q_t(s, a) \leftarrow Q_t(s, a) + \alpha_n(s, a)(r + \gamma \max_{a' \in A} Q_{t-1}(s', a') - Q_t(s, a))$
3 **if** $t = T$ **then**
 | $s \leftarrow s_0$
 | $t \leftarrow 1$
 else
 | $s \leftarrow s'$
 | $t \leftarrow t + 1$
 end
end
end

Algorithm 1: Q-learning

optimal policy, which is rarely true in practice. In interactive settings (Regan and Boutilier, 2009; Weng and Zanuttini, 2013), this elicitation process can be cognitively very complex as it requires to balance several criteria in a complex manner and as it can imply a large number of parameters. In this paper, we address this problem by only assuming that a strict weak ordering over histories is known.

Secondly, for numerous applications, the expectation of cumulated reward, as used in Equation 1, may not be the most appropriate criterion (even when a numeric reward function is defined). For instance, in case of high variance or when a policy is known to be only applied a few times, the solution given by this criterion may not be satisfying for risk-averse agent. Moreover, in some domains (*e.g.*, web industry or more generally service industry), decisions about performance are often based on the minimal quality of 99% of the possible outcomes. Therefore, in this article we aim at using a quantile (defined in Section 3.5) as a decision criterion to solve an MDP.

3.4. MDP with End States

In this paper, we work with episodic MDPs with end states. Such an MDP is formally defined as a tuple $\mathcal{M}_T = (S, A, G, \mathbf{P}, s_0)$ where S, A, \mathbf{P}, s_0 are defined as previously, $G \subseteq S$ is a finite set of end states and T is a finite maximal horizon (*i.e.*, an end state is attained after at most T time steps.). We call *episode* a history starting from s_0 and ending in a final state of G .

We assume that a preference relation is defined over end states: We write $g' \prec g$ if end state g is preferred to end state g' . Without loss of generality, we assume that $G = \{g_1, \dots, g_n\}$ and end states are ordered with increasing preference, *i.e.*, $g_1 \prec g_2 \prec \dots \prec g_n$. The weak relation of \prec is denoted \preceq .

Note that a finite horizon MDP can be reformulated as an MDP with end states by state augmentation. Although the resulting MDP may have a large-sized state space, the two models are formally equivalent. We focus on episodic MDPs with end states to simplify the presentation of our approach.

3.5. Quantile Criterion

We define quantiles of distributions over end states of G , which are ordered by \preceq . Let $\tau \in [0, 1]$ be a fixed parameter. Intuitively, the τ -quantile of a distribution of end states, is the value $q \in G$ such that the probability of getting an end state equal or lower than q is τ and that of getting an end state equal or greater than q is $1 - \tau$. The 0.5-quantile, also known as median, can be seen as the ordinal counterpart of the mean. The 0-quantile (resp. 1-quantile) is the minimum (resp. maximum) of a distribution. More generally, quantiles, which have been axiomatically characterized by Rostek (2010), define decision criteria that have the nice property of not requiring numeric valuations, but only an order.

The formal definition of quantiles can be stated as follows. Let p^π denote the probability distribution over end states induced by a policy π from initial state s_0 , the *cumulative distribution* induced by p^π is then defined as F^π where $F^\pi(g) = \sum_{g' \preceq g} p^\pi(g')$ is the probability of getting an end state not preferred to g when applying policy π . Similarly, the *decumulative distribution* induced by p^π is defined as $G^\pi(g) = \sum_{g \preceq g'} p^\pi(g')$ is the probability of getting an end state not lower than g .

These two notions of cumulative and decumulative enable us to define two kinds of criteria. First, given a policy π , we define the lower τ -quantile for $\tau \in (0, 1]$ as:

$$\underline{q}_\tau^\pi = \min\{g \in G \mid F^\pi(g) \geq \tau\} \tag{2}$$

where the min operator is with respect to \preceq .

Then, given a policy π , we define the upper τ -quantile for $\tau \in [0, 1)$ as:

$$\bar{q}_\tau^\pi = \max\{g \in G \mid G^\pi(g) \geq 1 - \tau\} \tag{3}$$

where the max operator is with respect to \preceq .

If $\tau = 0$ or $\tau = 1$ only one of \underline{q}_τ^π or \bar{q}_τ^π is defined and we define the τ -quantile q_τ^π as that value. When both are defined, by construction, we have $\underline{q}_\tau^\pi \preceq \bar{q}_\tau^\pi$. If those two values are equal, q_τ^π is defined as equal to them. For instance, this is always the case in continuous settings for continuous distributions. However, in our discrete setting, it could happen that those values differ, as shown by Example 1.

Example 1 Consider an MDP where $G = \{g_1 \prec g_2 \prec g_3\}$. Let π be a policy that attains each end state with probabilities 0.5, 0.2 and 0.3 respectively. It is easy to check that $\underline{q}_{0.5}^\pi = g_1$ whereas $\bar{q}_{0.5}^\pi = g_2$.

When the lower and the upper quantiles differ, one may define the quantile as a function of the lower and upper quantiles (Weng, 2012). For simplicity, in this paper, we focus on optimizing directly the lower and the upper quantiles.

The quantile criterion is difficult to optimize, even when a numerical reward function is given and the quality of an episode is defined as the cumulative of rewards received along the episode. This difficulty comes notably from two related sources:

- The quantile criterion is *non-linear*: for instance, the τ -quantile $q_\tau^{\tilde{\pi}}$ of the mixed policy $\tilde{\pi}$ that generates an episode using policy π with probability p and π' with probability $1 - p$ is not equal to $pq_\tau^\pi + (1 - p)q_\tau^{\pi'}$.
- The quantile criterion is *non-dynamically consistent*: A sub-policy at time step t of an optimal policy for horizon T may not be optimal for horizon $T - t$.

In decision theory (McClennen, 1990), three approaches have been considered for such kinds of decision criteria:

1. *Consequentialist* approach: at each time step t , follow an optimal policy for the problem with horizon $T - t$ and initial state s_t even if the resulting policy is not optimal at horizon T ;
2. *Resolute choice* approach: at time step $t = 0$, apply an optimal policy for the problem with horizon T and initial state s_0 and do not deviate from it;
3. *Sophisticated resolute choice* approach (Jaffray, 1998; Fargier et al., 2011): apply a policy π (chosen at the beginning) that trades off between how much π is optimal for all horizons $T, T - 1, \dots, 1$.

With non-dynamically consistent preferences, it is debatable to adopt a consequentialist approach, as the sequence of decisions may lead to dominated results. In this paper, we adopt a resolute choice point of view. We leave the third approach for future work.

4. Quantile-based Reinforcement Learning

In this section, we first state the problem solved in this paper and some useful properties. Then, we present our algorithm called Quantile Q-learning (or QQ-learning for short), which is an extension of Q-learning and exploits a two-timescale stochastic approximation technique.

4.1. Problem Statement

In this paper, we aim at learning a policy that is optimal for the quantile criterion from a fixed initial state. We assume that the underlying MDP is an episodic MDP with end states. Let $\tau \in (0, 1)$ be a fixed parameter. Formally, the problem of determining a policy optimal for the lower/upper τ -quantile can be stated as follows:

$$\pi^* = \arg \max_{\pi} \underline{q}_\tau^\pi \quad \text{or} \quad \pi^* = \arg \max_{\pi} \bar{q}_\tau^\pi \quad (4)$$

We focus on learning a policy that is deterministic and Markovian.

The optimal lower/upper quantiles satisfy the following lemmas:

Lemma 1 *The optimal lower τ -quantile \underline{q}_τ^* satisfies:*

$$\underline{q}_\tau^* = \min\{g : F^*(g) \geq \tau\} \quad (5)$$

$$F^*(g) = \min_{\pi} F^\pi(g) \quad \forall g \in G \quad (6)$$

and the optimal upper τ -quantile \bar{q}_τ^* satisfies:

$$\bar{q}_\tau^* = \max\{g : G^*(g) \geq 1 - \tau\} \quad (7)$$

$$G^*(g) = \max_{\pi} G^\pi(g) \quad \forall g \in G \quad (8)$$

Then, if the optimal lower quantile (q_τ^*) or upper quantile (\bar{q}_τ^*) were known, the problem would be relatively easy to solve. By Lemma 1, an optimal policy for the lower quantile could be obtained as follows:

$$\pi^* = \arg \min_{\pi} F^\pi(\underline{q}_\tau^-) \quad (9)$$

where $\underline{q}_\tau^- = g_1$ if $q_\tau^* = g_1$ and $\underline{q}_\tau^- = g_i$ if $q_\tau^* = g_{i+1}$. The reason one needs to use \underline{q}_τ^- instead of the optimal lower quantile is that otherwise it may happen that the cumulative distribution of a non-optimal policy π is smaller than or equal to that of an optimal policy at the lower quantile q_τ^* (but greater below \underline{q}_τ^- as π is not optimal). In that case, the lower quantile of π might not be q_τ^* . Such a thing cannot happen for the upper quantile. In that sense, determining an optimal policy for the upper quantile is easier. By Lemma 1, it is simply given by:

$$\pi^* = \arg \max_{\pi} G^\pi(\bar{q}_\tau^*) \quad (10)$$

In practice, if the lower/upper quantiles were known, those policies could be computed by solving a standard MDP with the following reward functions:

$$\underline{\mathbf{R}}_\theta(s) = \begin{cases} 0 & \forall s \notin \mathcal{G} \\ 0 & \forall s = g_i \in \mathcal{G}, \theta \geq i \\ 1 & \forall s = g_i \in \mathcal{G}, \theta \leq i - 1 \end{cases}$$

for the lower quantile with $\theta = k$ if $\underline{q}_\tau^- = g_k$ and

$$\bar{\mathbf{R}}_\theta(s) = \begin{cases} 0 & \forall s \notin \mathcal{G} \\ 1 & \forall s = g_i \in \mathcal{G}, \theta \leq i \\ 0 & \forall s = g_i \in \mathcal{G}, \theta \geq i + 1 \end{cases}$$

for the upper quantile with $\theta = k$ if $\bar{q}_\tau^* = g_k$. Note that $\underline{\mathbf{R}}_\theta$ can be rewritten to depend on the optimal lower quantile:

$$\underline{\mathbf{R}}_\theta(s) = \begin{cases} 0 & \forall s \notin \mathcal{G} \\ 0 & \forall s = g_i \in \mathcal{G}, \theta \geq i + 1 \\ 1 & \forall s = g_i \in \mathcal{G}, \theta \leq i \end{cases}$$

with $\theta = k$ if $q_\tau^* = g_k$.

Solving an MDP with $\underline{\mathbf{R}}_\theta$ amounts to minimizing the probability of ending in a final state strictly less preferred than q_τ^* , which solves Equation 9. Similarly, solving an MDP with $\bar{\mathbf{R}}_\theta$ amounts to maximizing the probability of ending in a final state at least as preferred as \bar{q}_τ^* , which solves Equation 10.

Now, the issue here is that the lower and upper quantiles are not known. We show in the next subsection that this problem can be overcome with a two-timescale stochastic approximation technique.

Data: $\mathcal{M}_T = (S, A, G, \mathbf{P}, s_0)$
begin
 Initialize θ to a random value
for $n = 1, 2, \dots$ **do**
 4 Solve MDP \mathcal{M}_T with reward function $\bar{\mathbf{R}}_\theta$
 if $V_\theta^*(s_0) < 1 - \tau$ **then**
 | $\theta \leftarrow \theta - 1/n$
 else
 | $\theta \leftarrow \theta + 1/n$
 end
end
end

Algorithm 2: Simple strategy for finding the optimal upper quantile

4.2. QQ-learning

As the lower and upper quantiles are not known, we let parameter θ vary in \mathbb{R}^+ during the learning steps and we refine the definition of the previous reward functions to make sure they are both well-defined for all $\theta \in \mathbb{R}^+$ and smooth in θ :

$$\underline{\mathbf{R}}_\theta(s) = \begin{cases} 0 & \forall s \notin \mathcal{G} \\ -1 & \forall s = g_i \in \mathcal{G}, \theta \geq i + 1 \\ 0 & \forall s = g_i \in \mathcal{G}, \theta \leq i \\ i - \theta & \text{else} \end{cases}$$

$$\bar{\mathbf{R}}_\theta(s) = \begin{cases} 0 & \forall s \notin \mathcal{G} \\ 1 & \forall s = g_i \in \mathcal{G}, \theta \leq i \\ 0 & \forall s = g_i \in \mathcal{G}, \theta \geq i + 1 \\ i + 1 - \theta & \text{else} \end{cases}$$

In the remaining of the paper, we present how to solve for the upper quantile. A similar approach can be developed for the lower quantile. In order to find the optimal upper quantile, one could use the strategy described in Algorithm 2. Value $V_\theta^*(s_0)$ approximates the probability of reaching an end state whose index is at least as high as θ . If that value is smaller than $1 - \tau$, it means θ is too high and should be decreased. Otherwise θ is too small and should be increased. Parameter θ will then converge to the index of the optimal upper quantile, which is the maximal value for θ such that $V_\theta^*(s_0) \geq 1 - \tau$. The optimal policy for V_θ^* is an optimal policy for the upper τ -quantile.

In a reinforcement learning setting, the solve MDP part (line 4 in Algorithm 2) could be replaced by an RL algorithm such as Q-learning. The problem is that such algorithm is only guaranteed to converge to the solution when $n \rightarrow \infty$. It would be therefore difficult to integrate Q-learning in Algorithm 2. Instead, a good policy can be learned while searching for the correct value of θ . To that aim, we use a two-timescale technique (Borkar, 1997, 2008) in which Q-learning and the update of parameter θ are run concurrently but at different speeds (*i.e.*, at two different timescales). For this to work, parameter θ needs to be seen as

Data: $\mathcal{M}_T = (S, A, G, \mathbf{P}, s_0)$ with \mathbf{P} unknown

Result: Q

begin

$Q(s, a) \leftarrow 0, \forall (s, a) \in S \times A$

$s \leftarrow s_0$

$t \leftarrow 1$

for $n = 1$ *to* N **do**

$a \leftarrow$ choose action

$r, s' \leftarrow$ perform action a in state s

$Q_t(s, a) \leftarrow Q_t(s, a) + \alpha_n(s, a)(r + \gamma \max_{a' \in A} Q_{t-1}(s', a') - Q_t(s, a))$

if $V_\theta^*(s_0) < 1 - \tau$ **then**

$\theta \leftarrow \theta - 1/n$

else

$\theta \leftarrow \theta + 1/n$

end

if $s' \in G$ **then**

$s \leftarrow s_0$

$t \leftarrow 1$

else

$s \leftarrow s'$

$t \leftarrow t + 1$

end

end

end

Algorithm 3: QQ-learning for the upper-quantile

quasi-static for the Q-learning algorithm. This is possible if the ratio of the learning rate of Q-learning and that of the update of θ satisfies:

$$\lim_{n \rightarrow \infty} \frac{\beta_n}{\alpha_n} = 0 \tag{11}$$

where $\beta_n = 1/n$ is the learning rate for parameter θ and α_n is the learning rate in the Q-learning algorithm. Equation 11 implies that parameter θ is changing at a slower timescale than the Q function.

5. Experimental Results

To demonstrate the soundness of our approach, we evaluate our algorithm on the domain, Who wants to be a millionaire. We present the experimental results below.

5.1. Domain

In this popular television game show, a contestant needs to answer a maximum of 15 multiple-choice questions (with four possible answers) of increasing difficulty, for increasingly large sums, roughly doubling the pot at each question. At each time step, the contestant may decide to walk away with the money currently won. If she answers incorrectly, then all winnings are lost except what has been earned at a “guarantee point” (questions

5 and 10). The player is allowed 3 lifelines (50:50, which removes two choices, ask the audience and call a friend for suggestions); each can only be used once. We used the first model of the Spanish 2003 version of the game presented by [Perea and Puerto \(2007\)](#). The probability of answering correctly is a function of the question’s number and increased by the lifelines used (if any).

5.2. Results

We plot the results (see [Figure 1](#)) obtained for two different learning runs, one with 1 million learning steps and the other with 10 million learning steps. The θ -updates were interleaved with a Q-learning phase using an ε -greedy exploration strategy with $\varepsilon = 0.01$ and with learning rate $\alpha(n) = 1/(n + 1)^{11/20}$. One can check that [Equation 11](#) is satisfied. We optimize the upper quantile with $\tau = 0.3$. During the learning process we maintain a vector \mathbf{f} of frequencies with which each final state has been attained. We define the quantity *score* as the cross product between \mathbf{f} and the vector of rewards obtained when attaining each final state given the current value of θ . Put another way, *score* is the value of the non-stationary policy that has been played since the beginning of the learning process. Moreover, at each iteration we compute $V_\theta^*(s_0)$, the optimal value in s_0 given the current value of θ .

On [Figures 1\(a\)](#) and [1\(b\)](#) we observe the evolution of $V_\theta^*(s_0)$ as the number of iterations increases. We observe that $V_\theta^*(s_0)$ converges towards $1 - \tau = 0.7$ as the number of iterations increases with oscillations of decreasing amplitude that are due to the ever changing θ value. [Figures 1\(c\)](#) and [1\(d\)](#) show the evolution of the *score* as the number of iterations increases. *Score* converges towards a value of 0.7 but inferior. This is due to the exploration of the Q-learning algorithm. Lastly, [1\(e\)](#) and [1\(f\)](#) plot the evolution of θ as the number of iteration increases. The value of θ converges towards a value of 4 which the one for which $V_\theta^*(s_0) = 1 - \tau = 0.7$.

6. Conclusion

We have presented an algorithm for learning a policy optimal for the quantile criterion in the reinforcement learning setting, when the MDP has a special structure, which corresponds to repeated episodic decision-making problems. It is based on stochastic approximation with two timescales ([Borkar, 2008](#)). Our proposition is experimentally validated on the domain, Who wants to be millionaire.

As future work, it would be interesting to investigate how to choose the learning rate α_n in order to ensure a fast convergence. Moreover, our approach could be extended to other settings than episodic MDPs. Besides, it would also be interesting to explore whether gradient-based algorithms could be developed for the optimization of quantiles, based on the fact that a quantile is solution of an optimization problem where the objective function is piecewise linear ([Koenker, 2005](#)).

QUANTILE RL

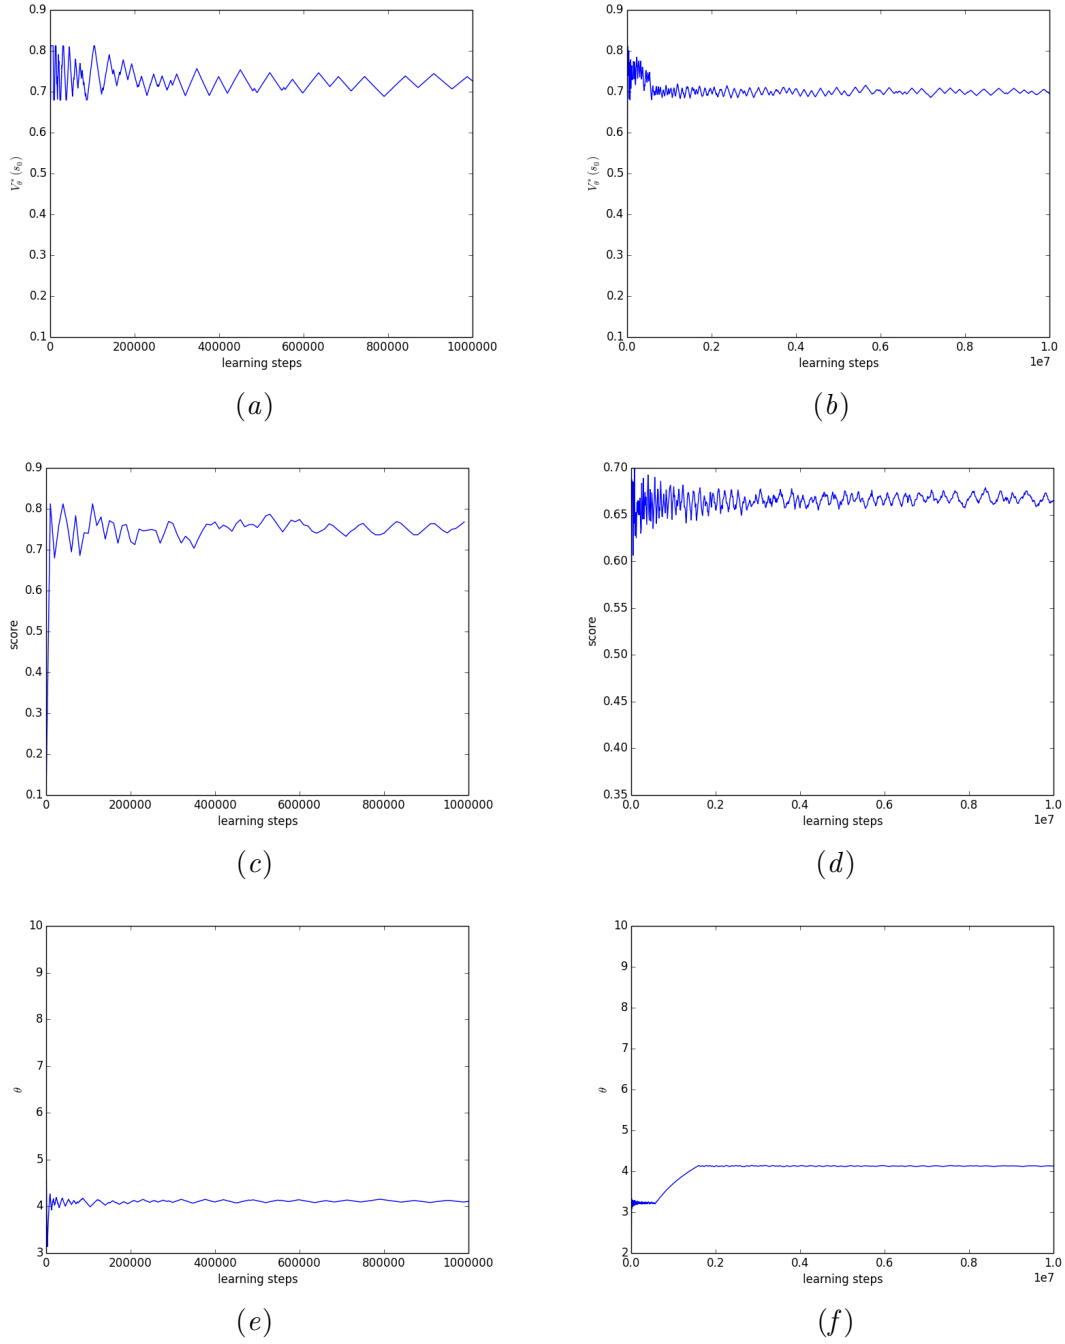


Figure 1: Evolution of $V_{\theta}^*(s_0)$, score and θ for 1 million learning steps (left) and 10 million learning steps (right)

References

- Pieter Abbeel, Adam Coates, and Andrew Y. Ng. Autonomous helicopter aerobatics through apprenticeship learning. *International Journal of Robotics Research*, 29(13):1608–1639, 2010.
- R. Akrou, M. Schoenauer, and M. Sébag. April: Active preference-learning based reinforcement learning. In *ECML PKDD, Lecture Notes in Computer Science*, volume 7524, pages 116–131, 2012.
- Nicole Bäuerle and Jonathan Ott. Markov decision processes with average value-at-risk criteria. *Mathematical Methods of Operations Research*, 74(3):361–379, 2011.
- D.F. Benoit and D. Van den Poel. Benefits of quantile regression for the analysis of customer lifetime value in a contractual setting: An application in financial services. *Expert Systems with Applications*, 36:10475–10484, 2009.
- D. Blackwell. An analog of the minimax theorem for vector payoffs. *Pacific Journal of Mathematics*, 6(1):1–8, 1956.
- V. Borkar and Rahul Jain. Risk-constrained Markov decision processes. *IEEE Transactions on Automatic Control*, 59(9):2574–2579, 2014.
- Vivek S. Borkar. *Stochastic approximation : a dynamical systems viewpoint*. Cambridge university press New Delhi, Cambridge, 2008. ISBN 978-0-521-51592-4. URL <http://opac.inria.fr/record=b1132816>.
- V.S. Borkar. Stochastic approximation with time scales. *Systems & Control Letters*, 29(5):291–294, 1997.
- M. Boussard, M. Bouzid, A.I. Mouaddib, R. Sabbadin, and P. Weng. *Markov Decision Processes in Artificial Intelligence*, chapter Non-Standard Criteria, pages 319–359. Wiley, 2010.
- Véronique Bruyère, Emmanuel Filiot, Mickael Randour, and Jean-François Raskin. Meet your expectations with guarantees: beyond worst-case synthesis in quantitative games. In *STACS*, 2014.
- R. Busa-Fekete, B. Szörenyi, P. Weng, W. Cheng, and E. Hüllermeier. Preference-based reinforcement learning. In *European Workshop on Reinforcement Learning, Dagstuhl Seminar*, 2013.
- Robert Busa-Fekete, Balazs Szorenyi, Paul Weng, Weiwei Cheng, and Eyke Hüllermeier. Preference-based Reinforcement Learning: Evolutionary Direct Policy Search using a Preference-based Racing Algorithm. *Machine Learning*, 97(3):327–351, 2014.
- Alexandra Carpentier and Michal Valko. Extreme bandits. In *NIPS*, 2014.
- Yinlam Chow and Mohammad Ghavamzadeh. Algorithms for cvar optimization in MDPs. In *NIPS*, 2014.
- G. DeCandia, D. Hastorun, M. Jampani, G. Kakulapati, A. Lakshman, A. Pilchin, S. Sivasubramanian, P. Voshall, and W. Vogels. Dynamo: amazon’s highly available key-value store. *ACM SIGOPS Operating Systems Review*, 41(6):205–220, 2007.
- E. Delage and S. Mannor. Percentile optimization in uncertain Markov decision processes with application to efficient exploration. In *ICML*, pages 225–232, 2007.

- Hélène Fargier, Gildas Jeantet, and Olivier Spanjaard. Resolute choice in sequential decision problems with multiple priors. In *IJCAI*, 2011.
- Jerzy A. Filar. Percentiles and Markovian decision processes. *Operations Research Letters*, 2(1): 13 – 15, 1983. ISSN 0167-6377. doi: [http://dx.doi.org/10.1016/0167-6377\(83\)90057-3](http://dx.doi.org/10.1016/0167-6377(83)90057-3). URL <http://www.sciencedirect.com/science/article/pii/0167637783900573>.
- Jerzy A. Filar, L. C. M. Kallenberg, and Huey-Miin Lee. Variance-penalized Markov decision processes. *Mathematics of Operations Research*, 14:147–161, 1989.
- P.C. Fishburn. An axiomatic characterization of skew-symmetric bilinear functionals, with applications to utility theory. *Economics Letters*, 8(4):311–313, 1981.
- J. Fürnkranz, E. Hüllermeier, W. Cheng, and S.H. Park. Preference-based reinforcement learning: A formal framework and a policy iteration algorithm. *Machine Learning*, 89(1):123–156, 2012.
- Hugo Gilbert, Olivier Spanjaard, Paolo Viappiani, and Paul Weng. Solving MDPs with skew symmetric bilinear utility functions. In *IJCAI*, pages 1989–1995, 2015.
- Hugo Gilbert, Bruno Zanuttini, Paolo Viappiani, Paul Weng, and Esther Nicart. Model-free reinforcement learning with skew-symmetric bilinear utilities. In *International Conference on Uncertainty in Artificial Intelligence (UAI)*, 2016.
- Hugo Gimbert. Pure stationary optimal strategies in Markov decision processes. In *STACS*, 2007.
- Jean-Yves Jaffray. Implementing resolute choice under uncertainty. In *UAI*, 1998.
- Philippe Jorion. *Value-at-Risk: The New Benchmark for Managing Financial Risk*. McGraw-Hill, 2006.
- D. Kalathil, V.S. Borkar, and R. Jain. A learning scheme for blackwell’s approachability in mdps and stackelberg stochastic games. *Arxiv*, 2014.
- R. Koenker. *Quantile Regression*. Cambridge university press, 2005.
- Y. Liu and S. Koenig. Risk-sensitive planning with one-switch utility functions: Value iteration. In *AAAI*, pages 993–999. AAAI, 2005.
- Y. Liu and S. Koenig. Functional value iteration for decision-theoretic planning with general utility functions. In *AAAI*, pages 1186–1193. AAAI, 2006.
- E. McClennen. *Rationality and dynamic choice: Foundational explorations*. Cambridge university press, 1990.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Belle-mare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518:529–533, 2015.
- A.Y. Ng and S. Russell. Algorithms for inverse reinforcement learning. In *ICML*. Morgan Kaufmann, 2000.
- F. Perea and J. Puerto. Dynamic programming analysis of the TV game who wants to be a millionaire? *European Journal of Operational Research*, 183(2):805 – 811, 2007. ISSN 0377-2217. doi: <http://dx.doi.org/10.1016/j.ejor.2006.10.041>. URL <http://www.sciencedirect.com/science/article/pii/S0377221706010320>.

- M.L. Puterman. *Markov decision processes: discrete stochastic dynamic programming*. Wiley, 1994.
- Mickaël Randour, Jean-François Raskin, and Ocan Sankur. Percentile queries in multi-dimensional Markov decision processes. *CoRR*, abs/1410.4801, 2014. URL <http://arxiv.org/abs/1410.4801>.
- K. Regan and C. Boutilier. Regret based reward elicitation for Markov decision processes. In *UAI*, pages 444–451. Morgan Kaufmann, 2009.
- M.J. Rostek. Quantile maximization in decision theory. *Review of Economic Studies*, 77(1):339–371, 2010.
- Balázs Szörényi, Róbert Busa-Fekete, Paul Weng, and Eyke Hüllermeier. Qualitative multi-armed bandits: A quantile-based approach. In *ICML*, pages 1660–1668, 2015.
- Gerald Tesauro. Temporal difference learning and td-gammon. *Communications of the ACM*, 38(3): 58–68, 1995.
- P. Weng. Markov decision processes with ordinal rewards: Reference point-based preferences. In *ICAPS*, volume 21, pages 282–289. AAAI, 2011.
- P. Weng. Ordinal decision models for Markov decision processes. In *ECAI*, volume 20, pages 828–833. IOS Press, 2012.
- P. Weng and B. Zanuttini. Interactive value iteration for Markov decision processes with unknown rewards. In *IJCAI*, 2013.
- D. J. White. Utility, probabilistic constraints, mean and variance of discounted rewards in Markov decision processes. *OR Spektrum*, 9:13–22, 1987.
- R. Wolski and J. Brevik. QPRED: Using quantile predictions to improve power usage for private clouds. Technical report, UCSB, 2014.
- Jia Yuan Yu and Evdokia Nikolova. Sample complexity of risk-averse bandit-arm selection. In *IJCAI*, 2013.
- Stella X. Yu, Yuanlie Lin, and Pingfan Yan. Optimization models for the first arrival target distribution function in discrete time. *Journal of mathematical analysis and applications*, 225:193–223, 1998.
- B. Zhang, Q. Cai, J. Mao, E Chang, and B. Guo. Spoken dialogue management as planning and acting under uncertainty. In *7th European Conference on Speech Communication and Technology*, 2001.